P⊕RTAL

US Patent & Trademark Office

Search: ⦿ The ACM Digital Library   ○ The Guide

(process <near> mode <near> thread) <near> ((cpu or proce

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

	Found
Terms used	**33,680**
process near mode near thread near cpu or **processor** or **processing near resource near allocate**	of
	134,837

Sort results by	relevance ▾	❖ Save results to a Binder	Try an Advanced Search
Display results	expanded form ▾	ⓘ Search Tips	Try this search in The ACM Guide
		☐ Open results in a new window	

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   next
Best 200 shown   Relevance scale ☐ ▭ ▨ ▨ ▨

1 Integrated document caching and prefetching in storage hierarchies based on Markov-chain predictions   ▨
Achim Kraiss, Gerhard Weikum
August 1998 **The VLDB Journal — The International Journal on Very Large Data Bases,**
Volume 7 Issue 3
Full text available: 📄 pdf(603.01 KB)   Additional Information: full citation, abstract, citings, index terms

Large multimedia document archives may hold a major fraction of their data in tertiary storage libraries for cost reasons. This paper develops an integrated approach to the vertical data migration between the tertiary, secondary, and primary storage in that it reconciles speculative prefetching, to mask the high latency of the tertiary storage, with the replacement policy of the document caches at the secondary and primary storage level, and also considers the interaction of these policies with ...

Keywords: Caching, Markov chains, Performance, Prefetching, Scheduling, Stochastic modeling, Tertiary storage

2 System-level power optimization: techniques and tools   ▨
Luca Benini, Giovanni de Micheli
April 2000 **ACM Transactions on Design Automation of Electronic Systems (TODAES),**
Volume 5 Issue 2
Full text available: 📄 pdf(385.22 KB)   Additional Information: full citation, abstract, references, citings, index terms

This tutorial surveys design methods for energy-efficient system-level design. We consider electronic sytems consisting of a hardware platform and software layers. We consider the three major constituents of hardware that consume energy, namely computation, communication, and storage units, and we review methods of reducing their energy consumption. We also study models for analyzing the energy cost of software, and methods for energy-efficient software design and compilation. This survery ...

3 Query evaluation techniques for large databases   ▨
Goetz Graefe
June 1993 **ACM Computing Surveys (CSUR),** Volume 25 Issue 2

Full text available: pdf(9.37 MB)   Additional Information: full citation, abstract, references, citings, index terms, review

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

4 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**

Full text available: pdf(4.21 MB)   Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

5 Piranha: a scalable architecture based on single-chip multiprocessing

Luiz André Barroso, Kourosh Gharachorloo, Robert McNamara, Andreas Nowatzyk, Shaz Qadeer, Barton Sano, Scott Smith, Robert Stets, Ben Verghese

May 2000 **ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture**, Volume 28 Issue 2

Full text available: pdf(191.10 KB)   Additional Information: full citation, abstract, references, citings, index terms

The microprocessor industry is currently struggling with higher development costs and longer design times that arise from exceedingly complex processors that are pushing the limits of instruction-level parallelism. Meanwhile, such designs are especially ill suited for important commercial applications, such as on-line transaction processing (OLTP), which suffer from large memory stall times and exhibit little instruction-level parallelism. Given that commercial applications constitute by fa ...

6 Experience Using Multiprocessor Systems—A Status Report

Anita K. Jones, Peter Schwarz

June 1980 **ACM Computing Surveys (CSUR)**, Volume 12 Issue 2

Full text available: pdf(4.48 MB)   Additional Information: full citation, references, citings, index terms

7 Performance characteristics of gang scheduling in multiprogrammed environments

Morris A. Jette

November 1997 **Proceedings of the 1997 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available: pdf(81.32 KB)   Additional Information: full citation, abstract, references, citings

Gang scheduling provides both space-slicing and time-slicing of computer resources for parallel programs. Each thread of execution from a parallel job is concurrently scheduled on an independent processor in order to achieve an optimal level of program performance. Time-slicing of parallel jobs provides for better overall system responsiveness and utilization than otherwise possible. Lawrence Livermore National Laboratory has deployed three generations of its gang scheduler on a variety of compu ...

Keywords: gang scheduling, multiprogramming, parallel system, scheduling, space-slicing, time-slicing

8 Concurrency control: methods, performance, and analysis
Alexander Thomasian
March 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 1

Full text available: pdf(427.18 KB)   Additional Information: full citation, references, citings, index terms

Keywords: Markov chains, adaptive methods, concurrency control, data contention, deadlocks, flow diagrams, load control, optimistic concurrency control, queueing network models, restart-oriented locking methods, serialiazability, thrashing, two-phase locking, two-phase processing, wait depth limited methods

9 Providing soft real-time QoS guarantees for Java threads
James C. Pang, Gholamali C. Shoja, Eric G. Manning
June 2001 **Proceedings of the 2001 joint ACM-ISCOPE conference on Java Grande**

Full text available: pdf(353.96 KB)   Additional Information: full citation, abstract, references, index terms

The Java platform has many characteristics that make it very desirable for integrated continuous media processing. Unfortunately, it lacks the necessary CPU resource management facility to support quality of service guarantees for soft real-time multimedia tasks. In this paper, we present our new Java Virtual Machine, Q-JVM, which brings CPU resource management to the Java platform. Q-JVM is based on Sun's JVM version 1.1.5. It implements an enhanced version or the MTR-LS algorithm in its thr ...

10 Models and languages for parallel computation
David B. Skillicorn, Domenico Talia
June 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 2

Full text available: pdf(298.05 KB)   Additional Information: full citation, abstract, references, citings, index terms

We survey parallel programming models and languages using six criteria to assess their suitability for realistic portable parallel programming. We argue that an ideal model should by easy to program, should have a software development methodology, should be architecture-independent, should be easy to understand, should guarantee performance, and should provide accurate information about the cost of programs. These criteria reflect our belief that developments in parallelism must be driven b ...

Keywords: general-purpose parallel computation, logic programming languages, object-oriented languages, parallel programming languages, parallel programming models, software development methods, taxonomy

11 Run-time adaptation in river
Remzi H. Arpaci-Dusseau

February 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 1

Full text available: pdf(849.04 KB)   Additional Information: full citation, abstract, references, index terms

We present the design, implementation, and evaluation of run-time adaptation within the River dataflow programming environment. The goal of the River system is to provide adaptive mechanisms that allow database query-processing applications to cope with performance variations that are common in cluster platforms. We describe the system and its basic mechanisms, and carefully evaluate those mechanisms and their effectiveness. In our analysis, we answer four previously unanswered and important que ...

Keywords: Performance availability, clusters, parallel I/O, performance faults, robust performance, run-time adaptation

12 The hierarchical simulation language HSL: a versatile tool for process-oriented simulation
D. P. Sanderson, R. Sharma, R. Rozin, S. Treu
April 1991 **ACM Transactions on Modeling and Computer Simulation (TOMACS)**, Volume 1 Issue 2

Full text available: pdf(2.68 MB)   Additional Information: full citation, references, citings, index terms, review

Keywords: C++, HSL, hierarchy, inheritance, interpreter, modularity, process, simulation programming language

13 Automatic compilation to a coarse-grained reconfigurable system-opn-chip
Girish Venkataramani, Walid Najjar, Fadi Kurdahi, Nader Bagherzadeh, Wim Bohm, Jeff Hammes
November 2003 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 2 Issue 4

Full text available: pdf(687.52 KB)   Additional Information: full citation, abstract, references, index terms

The rapid growth of device densities on silicon has made it feasible to deploy reconfigurable hardware as a highly parallel computing platform. However, one of the obstacles to the wider acceptance of this technology is its programmability. The application needs to be programmed in hardware description languages or an assembly equivalent, whereas most application programmers are used to the algorithmic programming paradigm. SA-C has been proposed as an expression-oriented language designed to im ...

Keywords: Reconfigurable computing, SIMD, compilers

14 Computing curricula 2001
September 2001 **Journal on Educational Resources in Computing (JERIC)**

Full text available: pdf(613.63 KB)   Additional Information: full citation, references, citings, index terms
html(2.78 KB)

15 Exokernel: an operating system architecture for application-level resource management
D. R. Engler, M. F. Kaashoek, J. O'Toole
December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles**, Volume 29 Issue 5

Full text available: 📄 pdf(2.16 MB)   Additional Information: full citation, references, citings, index terms

16 Virtual machine monitors: Xen and the art of virtualization

Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer,
Ian Pratt, Andrew Warfield
October 2003 **Proceedings of the nineteenth ACM symposium on Operating systems
                principles**

Full text available: 📄 pdf(168.76 KB)   Additional Information: full citation, abstract, references, index terms

Numerous systems have been designed which use virtualization to subdivide the ample
resources of a modern computer. Some require specialized hardware, or cannot support
commodity operating systems. Some target 100% binary compatibility at the expense of
performance. Others sacrifice security or functionality for speed. Few offer resource
isolation or performance guarantees; most provide only best-effort provisioning, risking
denial of service.This paper presents Xen, an x86 virtual machine monit ...

Keywords: hypervisors, paravirtualization, virtual machine monitors

17 An analysis of operating system behavior on a simultaneous multithreaded architecture

Joshua A. Redstone, Susan J. Eggers, Henry M. Levy
November 2000 **Proceedings of the ninth international conference on Architectural
                support for programming languages and operating systems**, Volume 28 ,
                34 Issue 5 , 5

Full text available: 📄 pdf(227.80 KB)   Additional Information: full citation, abstract, references, citings, index
                                                                  terms

This paper presents the first analysis of operating system execution on a simultaneous
multithreaded (SMT) processor. While SMT has been studied extensively over the past 6
years, previous research has focused entirely on user-mode execution. However, many of
the applications most amenable to multithreading technologies spend a significant fraction
of their time in kernel code. A full understanding of the behavior of such workloads
therefore requires execution and measurement of the operating sy ...

18 An analysis of operating system behavior on a simultaneous multithreaded architecture

Joshua A. Redstone, Susan J. Eggers, Henry M. Levy
November 2000 **ACM SIGPLAN Notices**, Volume 35 Issue 11

Full text available: 📄 pdf(1.56 MB)   Additional Information: full citation, abstract, references, index terms

This paper presents the first analysis of operating system execution on a simultaneous
multithreaded (SMT) processor. While SMT has been studied extensively over the past 6
years, previous research has focused entirely on user-mode execution. However, many of
the applications most amenable to multithreading technologies spend a significant fraction
of their time in kernel code. A full understanding of the behavior of such workloads
therefore requires execution and measurement of the operating sy ...

19 On multiprocessor system scheduling

Xiaotie Deng, Patrick Dymond
June 1996 **Proceedings of the eighth annual ACM symposium on Parallel algorithms
              and architectures**

Full text available: 📄 pdf(823.27 KB)   Additional Information: full citation, references, citings, index terms

20 Programming languages for distributed computing systems

Henri E. Bal, Jennifer G. Steiner, Andrew S. Tanenbaum
September 1989 **ACM Computing Surveys (CSUR)**, Volume 21 Issue 3

Full text available: pdf(6.50 MB)   Additional Information: full citation, abstract, references, citings, index terms, review

When distributed systems first appeared, they were programmed in traditional sequential languages, usually with the addition of a few library procedures for sending and receiving messages. As distributed applications became more commonplace and more sophisticated, this ad hoc approach became less satisfactory. Researchers all over the world began designing new programming languages specifically for implementing distributed applications. These languages and their history, their underlying pr ...

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10 next